

Google's Green Data Centers: Network POP Case Study

Table of Contents

Introduction	2
Best practices: Measuring performance, optimizing air flow, and turning up the thermostat	2
Best Practice #1 – Measuring performance	2
Best Practice #2 – Optimizing air flow	3
Best Practice #3 – Turning up the thermostat	3
Introducing the POP	3
Immediate improvements	4
Cold aisle containment	6
CRAC air return extensions	6
Adding a central CRAC controller	9
Final outcome: Energy savings and ROI	9
Energy savings	9
ROI analysis	10

Introduction

Every year, Google saves millions of dollars and avoids emitting tens of thousands of tons of carbon dioxide thanks to our data center sustainability efforts. In fact, our facilities use half the energy of a typical data center. This case study is intended to show you how you can apply some of the cost-saving measures we employ at Google to your own data centers and networking rooms.

At Google, we run many large proprietary data centers, but we also maintain several smaller networking rooms, called POPs or “Points of Presence”. POPs are similar to millions of small and medium-sized data centers around the world. This case study describes the retrofit of one of these smaller rooms, describing best practices and simple changes that you can make to save thousands of dollars each year.

For this retrofit, Google spent a total of \$25,000 to optimize this room’s airflow and reduce air conditioner use. A \$25,000 investment in plastic curtains, air return extensions, and a new air conditioner controller returned a savings of \$67,000/year. This retrofit was performed without any operational downtime.

Best practices: Measuring performance, optimizing air flow, and turning up the thermostat

Best Practice #1 – Measuring performance

The first step to managing efficiency in a POP or data center is to continuously measure energy usage focusing on two values:

- **IT equipment energy:** IT equipment energy is the energy consumed by servers, storage and networking devices—the machines that perform IT work.
- **Facility overhead energy:** The facility overhead energy is the energy used by everything else, including power distribution, cooling and lighting.

Power Usage Efficiency, or PUE, is the measurement used to compare these two types of energy. In other words, PUE is a measure for how efficiently a building delivers energy to the IT equipment inside. The ideal PUE is 1.0, meaning there’s no facility overhead energy—every watt of power going into the building is going straight to the computers and nowhere else.

$$PUE = \frac{IT\ Equipment\ Energy + Facility\ Overhead\ Energy}{IT\ Equipment\ Energy}$$

PUE must be measured over a long period of time for it to prove useful. At Google we look at both quarterly and trailing twelve-month performances. Snapshots of only a few hours are not helpful to make meaningful reductions in energy use.

Best practice #2 – Optimizing air flow

In a typical data center, the IT equipment is organized into rows, usually with a “cold aisle” in front where cold air enters the equipment racks and a “hot aisle” in back where hot air is exhausted. Computer room air conditioners, called CRACs, push cold air into the cold aisle, which flows through computer and network equipment into the hot aisle, where it returns to the CRAC. Cooling is the largest contributor to facility overhead energy.

The most important step in optimizing air flow is preventing hot and cold air from mixing. There is no single right way to do this. Being creative and finding simple ways to block and redirect air can greatly reduce the amount of cooling required. This includes simple things like installing blanking panels in empty rack slots and tightly sealing gaps in and around machine rows. It’s very similar to weatherizing your home.

It’s also important to eliminate any hot spots in order to achieve a more uniform thermal ‘profile.’ Localized hot spots are problematic for machines and trigger CRACs to turn on unnecessarily. Proper placement of temperature monitors and using computer modeling helps to quickly identify and eliminate hot spots.

Best practice #3 – Turning up the thermostat

It has long been believed that IT equipment needs to run at low temperatures—between 15°C/60°F and 21°C/70°F. However, the American Society of Heating, Refrigerating and Air Conditioning Engineers (ASHRAE) recommends cold aisle temperatures of up to 27°C/81°F, which we’ve found to have no detrimental effect on equipment. Most IT equipment manufacturers spec machines at 32°C/90°F or higher, so there is plenty of margin. In addition, most CRACs are set to dehumidify the air down to 40% relative humidity and to reheat air if the return air is too cold. Raising the temperature and turning off dehumidifying and reheating provides significant energy savings.

An elevated cold aisle temperature allows CRACs to operate more efficiently at higher intake temperatures. Also, it allows for more days of “free cooling”—days where mechanical cooling doesn’t need to run—if the facility has air- or water-side economization.

The simple act of raising the temperature from 22°C/72°F to 27°C/81°F in a single 200kW networking room could save tens of thousands of dollars annually in energy costs.

Introducing the POP

When retrofitting our POP, we had to work with what was already in the room. This meant no large-scale capital improvements and the room had to stay operational throughout the retrofit. We had to get creative and focus on maximizing the efficiency of the equipment that was already in place. With a few minor tweaks and some careful measurements, we were able to find significant savings.

Our starting point was a typical Tier III+ computing room with the following configuration:

- **Power:** Double-conversion uninterruptible power supplies (UPSs) with a designed IT load of 250kW.
- **Cooling:** Four 111kW computer room air conditioners (CRACs) with direct expansion evaporative cooling coil and remote air cooled condensers.
- **Racks:** Populated with commercially available third-party equipment, including optical switches, network routers, power supplies and load balancers.



A typical Google POP

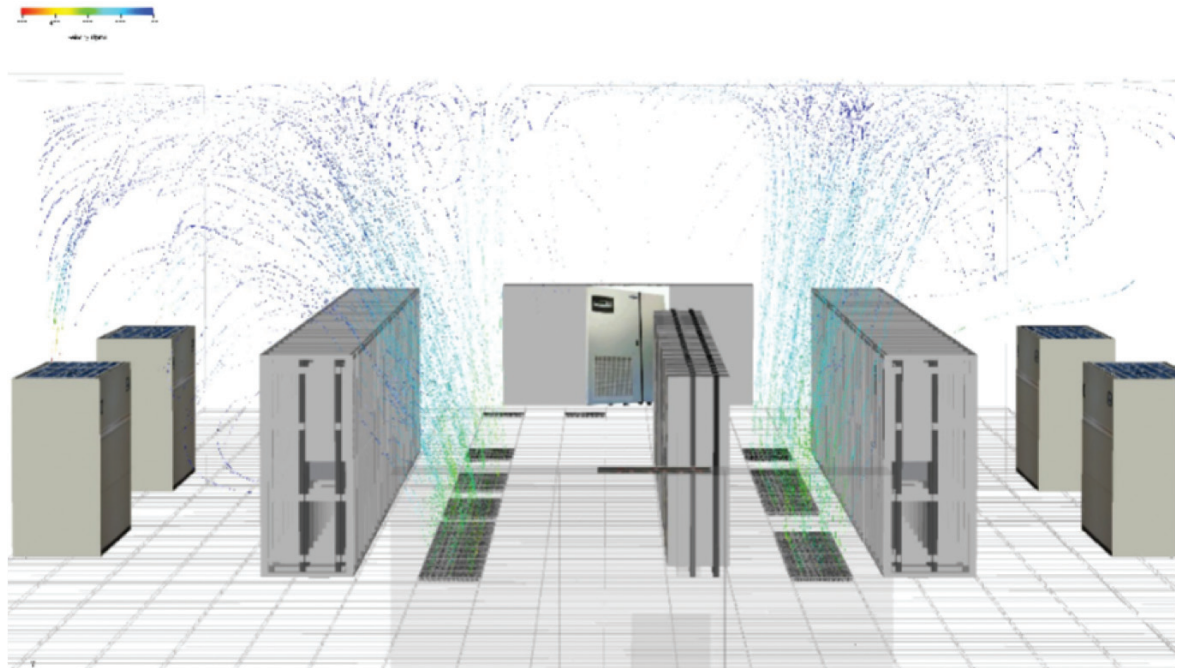
We looked at the settings on the CRACs and noticed their thermostats were set to 22°C/71°F at 40% relative humidity. The IT load for the room was only 85kW at the time, yet it was designed to hold 250kW of computer and network equipment. There were no attempts to optimize air flow. In this configuration, the room was overpowered, overcooled and underused.

When we took an initial measurement of PUE, **we found a high PUE of 2.4.**¹

Immediate improvements

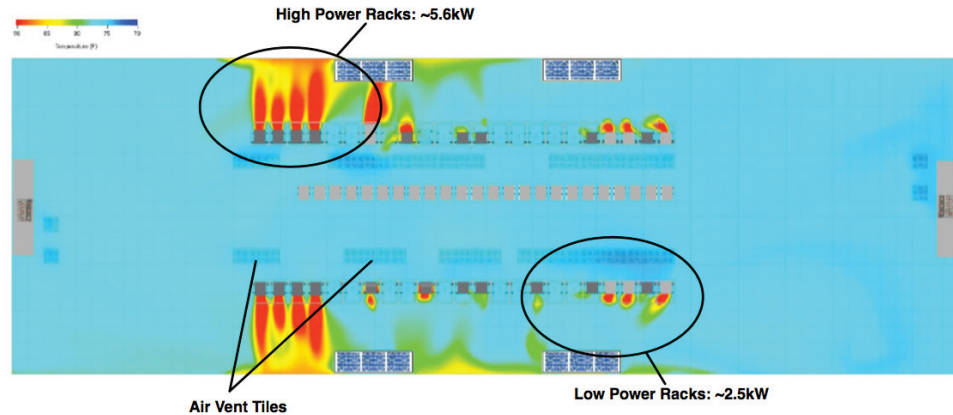
We first made simple changes before moving on to larger-scale improvements. In addition to PUE, we needed to measure cooling. Our hope was to improve cooling to the point that we could shut off a few CRACs.

After installing temperature monitors and examining airflow, we created thermal models using computational fluid dynamics (CFD) to run airflow simulations. As you can see from our computer model, most of the cold air bypasses the machines, coming up through the floor vents in the cold aisle, going over the machines, and plunging straight into the hot aisle.



Baseline model of airflow

Temperature sensors also identified hot spots in the POP. Racks with high-power machines consumed about 5.6kW per rack and produced more heat than racks with low-power or less densely populated machines. Since all the high-power racks were located in one area, that side of the room was significantly hotter than the other side of the room.



To optimize cooling and reduce energy usage, we took these immediate steps:

Step 1: Established critical monitoring points (CMPs)

In order to effectively measure temperature changes in our POP, we determined critical monitoring points (CMPs) where we needed accurate readings of temperature. These points included the room's UPS system, which specified a temperature below 25°C/77°F, and the input and output temperatures from the room's air conditioners.

Step 2: Optimized air vent tiles

We moved air vent tiles from the side with low power racks to the side with high power racks, matching airflow with IT power and significantly reducing the large hot spot on that side. Using our critical monitoring points, we were able to find the most efficient arrangement of vent tiles.

Step 3: Increased temperature and relative humidity settings

The original settings on the CRACs were a dry and chilly 22°C/71°F at a relative humidity of 40%. We implemented the latest ASHRAE recommended temperature range of 27°C/81°F max and increased the recommended humidity range to accept 20%-80% relative humidity.

Step 4: Contained UPS

The UPS required a temperature less than 25°C/77°F. However, the cold aisle air from the CRACs was now allowed to be a few degrees warmer than that. To prevent our UPS from getting too warm, we set up containment curtains (NFPA 701 compliant) to isolate the UPS. These containment curtains are similar to the plastic curtains used in commercial refrigerators.



Protected UPS

Step 5: Improved CRAC unit controller

Since we couldn't replace the air conditioners entirely, we decided to adjust the CRAC unit controllers to act smarter. We changed the settings on the CRACs' controllers to decrease their sensitivity to temperature changes and relative humidity, preventing the CRACs from turning on unnecessarily. We also disabled dehumidification and reheating functions and increased the amount of time required to activate cooling.

After implementing these changes, **we were happy to see a drop in PUE from 2.4 to 2.2.**

Cold aisle containment

Our original goal was to try to make our cooling efficient enough to shut off a few CRACs and save energy. In order to get there, we'd have to increase the efficiency of cold air entering the cold aisle air by raising the temperature difference between the cold and hot aisle so that only the hottest air is being returned to the CRACs.

To increase our cold aisle efficiency, we investigated a few ways to seal off the cold aisles from the hot aisles. Our first design involved creating a full containment unit that included putting a lid over the cold aisle. Unfortunately, that involved modifying the sprinkler system in order to comply with local fire codes.

Instead, we sealed the cold aisles by using blanking plates on the backs of empty rack space and adding refrigerator door curtains to the ends of cold aisles. We also added refrigerator curtains every 3 meters along the length of the cold aisle to better direct air flow.



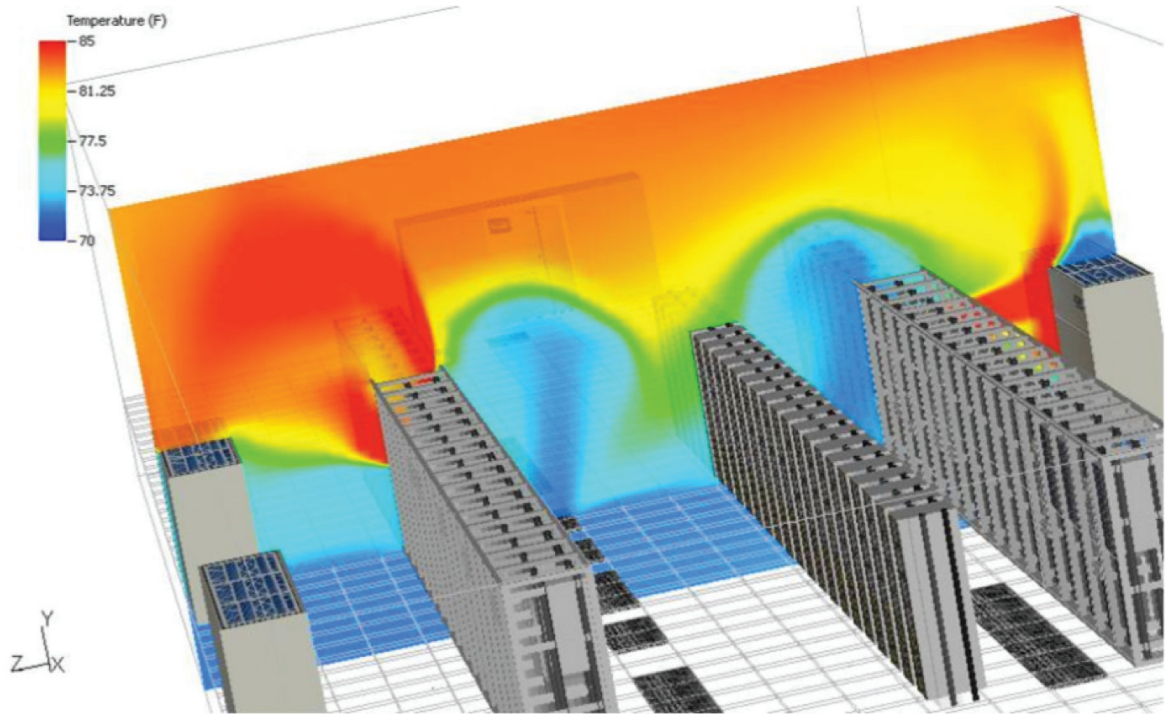
Cold aisle refrigerator curtains

After sealing off the cold aisle, **we saw another 0.2 drop in PUE to 2.0.**

CRAC air return extensions

A CFD simulation revealed that we had two specific issues with the CRAC hot air return:

- Hot air from a densely populated rack flowed directly into the CRAC air return, giving a falsely elevated temperature reading of the hot aisle. The false reading of an elevated temperature energized this CRAC more often than necessary.
- At a different CRAC, cold and hot air mixed directly at the CRAC air return, again reducing CRAC efficiency.



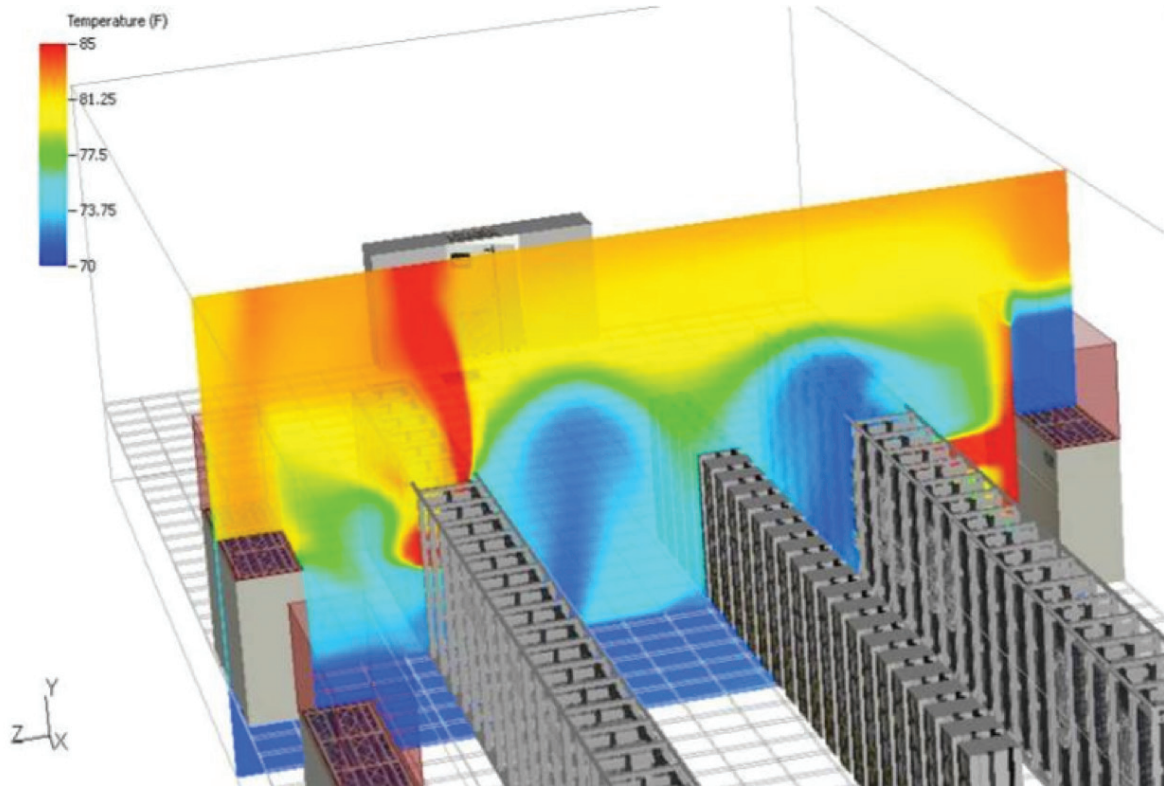
Hot aisle hot spot on left side

The simplest solution was to add sheet metal boxes that increased the height of the air returns by 1.2m/48in, leading to improvements in return air flow and the creation of a more uniform return air temperature to the CRAC.



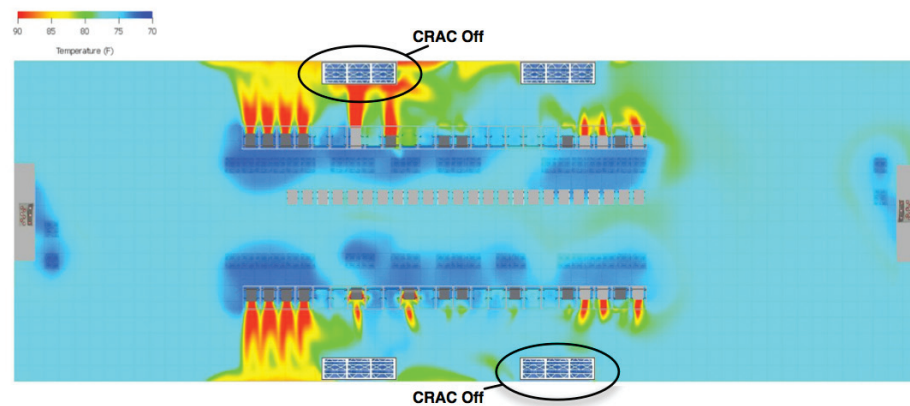
CRAC with an air return extension

By evening out the temperature in the hot aisle, we also prevented hot spots from building up. In particular, a hot spot above one of the high power racks was triggering a temperature sensor to turn on the compressor unnecessarily.



Hot aisle after adding air return extensions

With these optimizations, we reduced the number of CRACs from 4 to 2 while maintaining our desired cold aisle temperature. If you compare our floor analysis at this point with our initial floor analysis, you'll notice that the cold aisle is cooler and the warm aisle is about the same—this is with the thermostat turned up and only half the number of air conditioners turned on.



Floor analysis with only 2 CRACs on

We also installed motion sensors to turn off the overhead lights when the room was unoccupied, which further reduced the electrical power. With these changes in place, **the PUE dropped from 2.0 to 1.7.**

Adding a central CRAC controller

We next considered ways to dynamically control the number of CRACs turned on at any given time. If we only needed one air conditioner, we'd just turn on one. If we needed three, we could turn on three. If one suddenly stopped working, we could automatically turn on another to replace it.

Setting this up required a central CRAC controller tied to the temperature monitors in the room. We purchased this controller from a third party and connected it to the CRACs using the existing building management system. Now, we could turn on the air conditioning based on the load of the room and still maintain the required 2N redundancy. Even though at 85kW the room remained underused, the central CRAC controller made the cooling of the room more energy proportional.

Installing this system increased efficiency, reduced maintenance and increased cooling system redundancy. Now we were only using as many CRACs as needed, usually between 1 and 2. With less use, each air conditioner lasts longer and requires less maintenance. The new setup also allows for redundancy—if one of the CRACs fails, another instantly turns on in its place.

We decided not to disable the local controllers on each of the CRACs. If the new CRAC controller we installed happens to fail, the air conditioning in the room would continue to operate on the old controllers. Taking a final PUE reading for our POP, **we had a PUE of 1.5.**

Final outcome: Energy savings and ROI

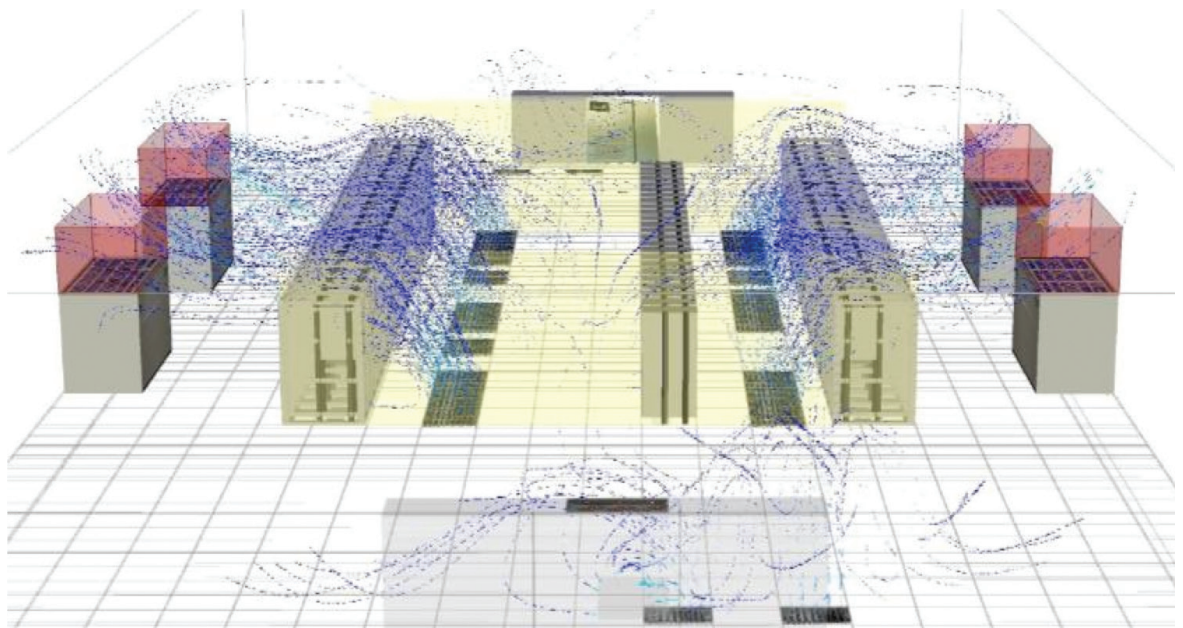
Energy savings

We've brought the facility overhead energy usage down from 1.4 to 0.5 per watt of IT computer energy, reducing it to a third of its original value. We managed to do this without any operational disruptions or causing any major temperature fluctuations within our POP.

The list of changes we made to the original space is fairly short:

1. Added temperature monitoring
2. Optimized air vent tiles
3. Increased temperature and relative humidity settings on CRACs
4. Blocked off the ends of cold aisles with curtains
5. Put blanking plates and side panels to block cold air passing through empty rack spaces
6. Added 48" extensions to all CRAC air returns
7. Added a new CRAC controller

Using computer modeling to analyze the airflow for our final setup, we can see a fairly efficient network room with far more air flowing through machines instead of bypassing them.



Final configuration airflow analysis

Having made changes to one POP, we applied the same changes to others, recording improvements each step of the way. PUE was checked using data collected every second over an 18 month period, and each PUE value contains an average of 86,400 data points. The results are consistent for every POP in which we implemented these efficiency improvements.

PUE improvements for five POPs

	POP 1	POP 2	POP 3	POP 4	POP 5
Starting point	2.4	2.2	2.2	2.4	2.4
After immediate improvements	2.2	2.0	2.0	2.2	2.0
After cold aisle containment	2.0	1.8	1.8	2.0	1.9
After adding CRAC air return extensions	1.7	1.7	1.7	1.7	1.7
After adding new CRAC controller	1.5	1.5	(Still collecting data)	1.6	(Still collecting data)

ROI analysis

Data center energy efficiency retrofits are a good example of where smart business and environmental stewardship coexist. In the POP in our case study, a capital investment of \$25,000 led to a yearly energy savings of over 670MWh, saving \$67,000 in yearly energy expenses. In addition, each improvement paid for itself in less than a year and will save hundreds of thousands of dollars throughout the lifetime of the equipment.

ROI for POP 1 improvements

For POP 1	PUE	Capital investment	PUE improvement	Savings/month	ROI (months)
Starting point	2.4	-	-	-	-
After immediate improvements	2.2	-	0.2	\$1,238	0
After cold aisle containment	2.0	12,000	0.2	\$1,238	9.7
After adding CRAC air return extensions	1.7	5,000	0.3	\$1,858	2.7
After adding new CRAC controller	1.5	8,000	0.2	\$1,238	6.5

The best practices described in this case study form the key elements of Google's power optimization strategies. Internet services have assumed a central role in daily life, driving demand for centralized networking and data centers. The energy demand created by this growth highlights the need for data center owners and operators to focus on power optimization as both an operational expense reduction and an environmental imperative.

1. Measuring PUE for our sites is described at: <http://www.google.com/corporate/datacenter/efficiency-measurements.html>.